

APPARATUS AND METHOD FOR REDUCING PRECISION OF DATA

5 CROSS-REFERENCE TO RELATED APPLICATION(S)

The present application claims the benefit of the filing date of United States Provisional Patent Application Serial No. 60/170,156, filed December 10, 1999, and entitled METHOD OF
10 ROUNDING, the contents of which are hereby expressly incorporated herein by reference.

15 BACKGROUND OF THE INVENTION

The implementation of DSP algorithms in VLSI requires trade-offs between factors such as algorithm performance, silicon area, power consumption, and clock frequency. One parameter that
20 affects all of these factors is the word length, or number of bits, used to represent the data. Where extremely precise computations are required, floating point arithmetic may be needed. However, floating point operations can be impractical
25 for high-speed applications, including digital signal processing (DSP), due to the added hardware, overhead, and processing time they demand.

By contrast, fixed point computations often are used when, such as in digital signal processing, extremely high processing rates are desirable. Typically, the word size used to represent a fixed point value is compact relative to floating point, and
30 fixed point computational devices tend to be of simpler,
35

1 37359/JFO/B600

5 efficient designs, which permit high-speed operation. Most DSP
algorithms achieve satisfactory performance in VLSI using fixed-
point arithmetic. Therefore, fixed point computations,
particularly in the two's complement binary format, are
predominant in DSP applications. However, because fixed point
10 systems can only represent a fixed range of numbers, the internal
data precision is reduced. In addition to the increased
possibility of overflow during a calculation, fixed point
operations tend to have a reduce ability to accurately quantize
15 small numbers. Also, fixed point operations tend to require
additional digits to represent the precision of the results of
an operation. For example, a fixed point multiplication of two
signal values, each having an initial precision of m bits,
20 results in a product having $2m$ bits. Because the word length of
a subsequent operation, or data path, may be limited to only m
bits, some form of precision reduction is needed.

25 Therefore, rounding is desirable where data processing
operations, including those involving, for example, binary- and
decimal-based operations, generate an undesirably large number
of digits of precision, and some form of precision reduction is
needed. Precision reduction in fixed point operations comes at
30 a price. For many DSP systems, performance metrics such as
signal-to-noise ratio (SNR) and stability are adversely affected
by decreasing internal data precision.

35

When the precision of a signal value is reduced, the difference between input and output values represents an error that is approximately equal to the part of the input signal that is discarded. In general, the statistics of the error can depend upon how the last bit of the output value is determined. Many DSP components, such as digital filters, can be represented by coefficient values related to a characteristic polynomial equation, which is representative of the component's operating characteristics, e.g. a digital filter's transfer function equation. Because the filter coefficients must be represented by finite length values, the behavior of the filter can be greatly influenced by precision reduction errors.

During the processing of a signal, and as the coefficients are updated, precision reduction can introduce errors roughly equivalent to the value of the dropped digits, or loss bits, thereby generating a precision reduction error signal. Signal processing typically involves numerous, sequential, iterative and recursive computations, during which uncompensated precision reduction errors accumulate, thereby degrading filter performance, possibly to the point of filter instability.

Many DSP applications include VLSI components that extensively employ high-speed computations which are susceptible to precision reduction errors. In widely-used adaptive filters and systems, the associated adaption components can be very sensitive to how data precision is reduced. In some instances,

1 37359/JFO/B600

5 precision reduction errors are merely nuisances with little practical impact. However, in the extreme, precision reduction errors can lead to disastrous outcomes, such as in the documented failure of a defensive missile battery during the 1991 Persian Gulf War, which resulted in many lost lives. Because DSP applications are pervasive in modern life, including for example, 10 communications, health care, transportation, defense, and the like, degraded system performance arising from uncompensated precision reduction errors can be of great import. Whether the application is aircraft navigation, critical-care life support, 15 electronic commerce, or global communications, it is imperative that the DSP infrastructure supporting the application operate in a reliable and robust manner to the greatest extent possible.

20 What is needed, then, are methods and apparatus that substantially eliminate or offset the precision reduction error.

25 SUMMARY OF THE INVENTION

The present invention satisfies the above needs by providing methods and apparatus for reducing precision of an input signal having a precision portion and a loss portion, by comparing the 30 loss portion to a preselected threshold value, f_t ; determining a selectable bias, α , responsive to the comparison of the loss portion, to the preselected threshold value, f_t ; and combining 35 the precision portion with α , thus creating a reduced precision

datum. Selectable bias α corresponds to a predetermined
characteristic of one of α , the input datum, the reduced
precision datum, and a combination thereof. In preferred
embodiments of the invention, selectable bias α is generated such
that the Expected Value of selectable bias α is substantially
equal to the Expected Value of the error signal between the input
data with higher precision and an output signal with lesser
precision. In this manner, errors due to rounding are minimized
or eliminated.

15 BRIEF DESCRIPTION OF THE DRAWINGS

These and other features, aspects and advantages of the present
invention will be more fully understood when considered with
respect to the following detailed description, appended claims
and accompanying drawings, wherein:

FIG. 1 is a signal flow representation of a precision
reduction processor;

FIG. 2A is a first depiction of general data structures
representative of the precision reduction;

FIG. 2B is a second, more specific depiction of data
structures representative of the precision reduction;

FIG. 3 is a first signal flow representation of one
embodiment of a selective bias rounding device, illustrating

general principles of operation according to the invention herein;

5 FIG. 4 is a second signal flow representation of another embodiment of a selective bias rounding device, illustrating more specific principles of operation according to the invention herein;

10 FIG. 5 is a first data flow diagram, illustrating one embodiment of the inventive methods herein;

15 FIG. 6 is a second data flow diagram, illustrating another embodiment of the inventive methods herein;

FIG. 7 is a third data flow diagram, illustrating yet another embodiment of the inventive methods herein;

20 FIG. 8 is a fourth data flow diagram, illustrating still another embodiment of the inventive methods herein;

FIG. 9 is a fifth data flow diagram, illustrating a further embodiment of the inventive methods herein;

25 FIG. 10 is a third signal flow representation of one embodiment of a selective bias rounding device, according to the present invention, illustrating particular structures therein;

30 FIG. 11 is a fourth signal flow representation of another embodiment of a selective bias rounding device, according to the present invention, illustrating memory-based control of a selectable bias;

35 FIG. 12 is a fifth signal flow representation of yet another embodiment of a selective bias rounding device, according to the

1 37359/JFO/B600

present invention, illustrating state-based control of a selectable bias;

5 FIG. 13 is a fifth signal flow representation of still another embodiment of a selective bias rounding device, according to the present invention, illustrating managed control of a selectable bias;

10 FIG. 14 is a sixth signal flow representation of still another embodiment of a selective bias rounding device, according to the present invention, illustrating adaptive parametric control of a selectable bias;

15 FIG. 15 is a VLSI system floor plan, illustrating a particular hardware implementation of managed control of a selectable bias;

20 FIG. 16 is a seventh signal flow representation of an SBR arithmetic unit, according to the present invention, in which a selectable bias rounding device is coupled to a standard arithmetic device; and

25 FIG. 17 is a generalized schematic of a 4-tap LMS adaptive filter, according to the present invention, having selectable bias rounding devices integrated therein for precision reduction management.

30 DETAILED DESCRIPTION OF THE INVENTION

The present invention includes methods and apparatus that
35 reduce the precision of an input signal value having a first

1 37359/JFO/B600

precision to an output signal having a second, lesser precision
in a manner that greatly reduces, or substantially cancels, a
5 precision reduction error signal typically inherent in prior art
rounding techniques. By combining the input signal with a
selectable bias, responsive to a preselected threshold rounding
10 state, the rounding methods and apparatus according to the
present invention provide an output signal that is substantially
free of precision reduction error bias. In addition, where it
is desired to produce a preselected signal offset, values for the
15 selectable bias can be assigned to generate the offset. Such a
signal offset may be useful to compensate for a undesirable pre-
existing input signal bias, including correcting for precision
reduction error biases injected during previous precision
20 reduction operations. Also, it may be useful to impart an
offset to an output signal, for example, to pre-condition the
output signal for an anticipated bias arising from subsequent
signal processing or in a communication channel.

25 As used herein, the term "selectable bias rounding" (SBR)
will be used with reference to the embodiments of the present
invention. SBR can be used alone, or in conjunction with
30 traditional hardware and software, to manage, or substantially
eliminate, the deleterious effects of accumulated round-off
error.

FIG. 1 illustrates a precision reduction processor 100,
35 which receives input signal Z 110. Error generator 140 produces

1 37359/JFO/B600

a precision reduction error signal 150 which, when combined with
input signal 110 in summer 130, creates reduced precision output
5 signal, \hat{Z} 120. Precision reduction error 150 can contain both
a component related to the extent of precision reduction and,
typically, a component related to the rounding technique used.
10 Symbolically, processor 120 can be represented by the equation: \hat{Z}
 $= Z + e$. It is clear that precision reduction error e is a
function of the difference between output signal \hat{Z} and input
signal Z . $e = \hat{Z} - Z$.

15 Therefore, SBR methods and apparatus that are targeted at
managing precision reduction error e can be advantageously used
to minimize, or substantially eliminate, certain forms of
precision reduction errors, thus allow DSP applications to
20 operate more robustly and reliably.

The present invention is not limited to digital signal
processing applications, or even to an electronic milieu.
Indeed, the SBR invention contemplates methods and apparatus that
25 may be used wherever it is desirable to manage error that may
result when reducing the precision of a datum, that is
representative of any physical entity. Also, although the SBR
methods and apparatus described herein are discussed in the
30 context of signed and unsigned decimal and binary data
representations, the present invention can be applied to other
data representation formats, and their complements, as will
35 become apparent to those of ordinary skill in the art. Moreover,

1 37359/JFO/B600

the present invention can be applicable to any precision
reduction errors, whether or not the errors are categorized as
5 "quantization," "truncation," or "rounding" errors, or a hybrid
thereof. Thus, SBR rounding methods and apparatus are useful to
minimize errors which may arise when, for example, a finite-
10 valued datum is transformed from a representation having greater
precision to one of lesser precision (e.g., 5.251 volts to 5
volts), or when a continuous-valued datum is approximated by a
corresponding a finite-valued datum (e.g., an analog datum to a
15 digital datum).

Finally, the principles of the present invention are
illustrated by exemplary signals (e.g., X , \hat{X} , Y , \hat{Y} , Z , and \hat{Z})
which may represent an individual signal sample; or a discrete
20 time sequence of signal samples, having a corresponding time
index. For example, input signal Z , and output signal \hat{Z} , in FIG.
1 can be described by the canonical form $Z(j)$ and $\hat{Z}(j)$, where j
is representative of the time index for a particular signal
25 datum. But, for simplicity, time indices such as j will be
dropped. A skilled artisan can discern when the discussion of
the exemplary signals relates to an individual signal sample, or
datum; or to a discrete time sequence of signal samples.
30 Although an SBR rounding operation can described with regard to
reducing the precision of an individual datum, or signal sample,
it is understood that references to statistical characteristics,
35 e.g., the signal Expected Value, or Mean, Variance, and the like,

1 37359/JFO/B600

are to be interpreted in the context of a stochastic process of
random variables, i.e., a signal that is defined by a time
5 sequence of signal samples that are random variables, of which
a particular datum is but a part.

Before elaborating upon the nature of selectable bias
10 rounding techniques and devices, it is useful to illustrate the
various concepts and terms related to "rounding," as used herein.
In general, there are five methods of rounding:

- (1) Round-to-zero (RTZ);
- 15 (2) Round-to-nearest (RTN);
- (3) Round-to-floor (RTF);
- (4) Round-to-ceiling (RTC); and
- (5) Round-to-even (RTE).

20 Each type of rounding tends to introduce some form of
precision reduction error. RTZ rounding tends to introduce a
cumulative downward error bias for positive signal values, and
a cumulative upper error bias for negative signal values, i.e.,
25 the magnitude of a signal value decreases. RTN rounding drives
the signal value to the nearest representable value. Although
the bias introduced in RTN rounding tends to be modest in
comparison with other modes of rounding, nevertheless significant
30 rounding errors can accumulate to produce undesirable results.
Using RTF rounding, both positive and negative signal values are
rounded towards negative infinity, introducing a negative
35 cumulative error bias into the signal. On the other hand, RTC

1 37359/JFO/B600

rounding rounds both positive and negative signal values towards
positive infinity, and a positive cumulative error bias is
5 introduced into the signal. RTC is often used when rounding
values which lie precisely half-way between two desired
quantities, e.g., rounding 3.50 to 4, or -3.50 to -3.) Finally,
10 RTE rounding involves rounding to the nearest even binary value.
A method employed by the floating point format specified by IEEE
Standard 754-1985 (International Standard IEC 559), the RTE
rounding mode adds the least significant bit (LSB) of the reduced
15 precision signal value itself, when the "half-way" value exists,
i.e., when the rounding digit following the LSB is non-zero, and
the extended precision result has a non-zero digit in any digit
location of the extended precision fields. As a result, the LSB
20 of the rounded value is always 0, an even value.

Truncation is an operation that is conceptually consistent
with RTZ rounding, in that, by dropping the undesired digits, the
resulting value is brought closer to zero. For example, dropping
25 the last two digits of -3.50 produces a value of -3. However, for
cases involving certain data representations, including negative
numbers in two's complement format, truncation effects the
opposite result, i.e., RTF rounding. For example, dropping the
30 last two digits of 100.10_2 (i.e., -3.5) produces 100_2 (i.e., -
3.0). Thus, truncation of two's complement signal values also
tends to introduce a negative cumulative error bias into the
35 signal. In view of this disparity, the term of "truncation" will

1 37359/JFO/B600

be replaced by the term for the appropriate rounding operation
wherever possible. Indeed, the term "precision reduction" is
5 generally used to include "truncation," "rounding," as well as
"quantization" or any other techniques that tend to reduce the
precision of a datum.

10 FIG. 2A depicts an exemplary conceptual framework for
illustrating a typical rounding operation. Input signal Y 200,
is stored in a rounding operand 210 during the precision
reduction operation. At the conclusion of the operation, operand
15 210 holds the value of output signal \hat{Y} . Input signal Y 200,
having $m+k$ digits, is received by rounding operand 210 for
processing. In this example, it is desired to reduce input signal
 Y 200 to output signal \hat{Y} 220 having m digits, by performing a
20 precision reduction operation on the value held by rounding
operand 210. As defined herein, the rounding operand includes
a precision portion 230, composed of m precision digits, followed
by a loss portion 240 composed of k loss digits. The m^{th}
25 precision digit is the least significant precision (LSP) digit
232 of precision portion 230, and the $(m+1)$ digit is rounding
digit 242, which is the most significant digit of the loss
portion 240. "Rounding point" 235, which may or may not
30 represent the radix point of the value rounding operand 210, lies
between (LSP) digit 232 and rounding digit 238. After the
completion of the rounding operation, loss portion 240 is
35 eliminated, leaving a signal having m precision digits. If

1 37359/JFO/B600

rounding point 235 coincides with the radix point for that value,
then the resultant reduced precision output signal value is an
5 integer; otherwise, the output signal value also contains a
fractional value and is real-valued. The value held in rounding
digit 235 can indicate that a threshold rounding state
10 potentially exists. One preferred threshold rounding state would
include the presence of a "half-value" in the loss portion.
Whether the threshold rounding state actually exists is
determined by examining the values of all k loss digits. For
15 example, in a decimal-oriented embodiment of the present
invention, the corresponding "half-value" in the rounding digit
235 may be 5_{10} . Under this scenario, if all digits subsequent to
the rounding digit 235 were zero-valued, then the threshold
20 rounding state would exist. Similarly, in a binary
implementation, if the rounding bit 235 holds a binary '1'
value, then the subsequent loss bits must be evaluated. If all
of the subsequent loss bits hold a zero value, then a threshold
25 rounding state exists. If any of the loss bits subsequent to the
rounding bit 235 is non-zero, then the threshold rounding state
does not exist. The existence of the threshold rounding state
can determine the type of rounding that will be employed to
30 generate Y . Typically, RTC rounding is used in prior art
precision reduction operations resulting from a threshold
rounding state. According to the present invention, it is
35 preferred to combine a selectable bias value α with the input

1 37359/JFO/B600

signal Y 200, to produce the reduced precision output signal Y 220, upon the occurrence of a threshold rounding state, e.g.,
5 when loss portion 240 exactly represents a "half-value".

FIG. 2B further illustrates the aforementioned rounding principles by depicting exemplary 8-bit binary input signal Y 250, rounding operand 260, and exemplary reduced precision 5-bit
10 binary output signal Y 270. In this illustration, bits b_7 through b_3 (251-255) represent the precision portion 256 of Y 250, and b_2 through b_0 (262-264) represent the loss portion 265.
15 Furthermore, the rounding point is between b_3 255, the LSP bit of the precision portion 256, and b_2 262 which is the rounding bit. If b_2 262 holds a non-zero value, for example, the potential for a threshold rounding state exists, and the values held in b_1 263
20 and b_0 264 are examined to verify the presence of the threshold rounding state. According to certain binary-oriented embodiments of the present invention, if rounding bit b_2 262 is "1,"
25 indicative of a "half-value," and if all of the subsequent bits, i.e., b_1 263 and b_0 264, are zero-valued, then the threshold rounding state exists. On the basis of loss portion 265, the desired method of rounding can be chosen and implemented.

30 In FIG. 3, selectable bias rounding method 300 encompasses a preselected precision reduction method in which a selectable bias α 310 is combined by summer 315 with input signal X 305, having a first precision, to produce output signal X 320 having
35 a second, lesser precision. SBR method 300 has a precision

reduction error signal, e 330, associated therewith, which could corrupt output signal X 320 except that values of bias α 310 are
 5 chosen to cancel the effect of error signal e 330. Error signal e 330 can include a component related to the extent to which the precision is being reduced (e.g., 24 bits signal reduced to 8
 10 bits signal), and a component due to the particular precision reduction technique that is employed (e.g., RTZ, RTN, RTC, RTF, RTE). Bias α 310 can be selected to cancel the effects if both components. For example, certain SBR techniques are designed to
 15 obviate the development of the latter bias, for example, by using alternating rounding. The value of the selectable bias α 310 is chosen to be responsive to a preselected characteristic of input signal X 305, the output signal X 320, the error signal e 330,
 20 values of α 310 itself, or combinations thereof. Selectable bias α 310 is preferred to be applied when input signal X 305 manifests a threshold rounding state, as represented by a preselected threshold value f_t 350. When the threshold rounding
 25 state does not exist, prior art rounding techniques may be employed.

For example, in one embodiment of the inventive method herein, it is preferred to choose values for selectable bias α
 30 310, corresponding to the Expected Values $E(e) = E(\alpha)$, when it is desired to substantially eliminate the precision reduction error signal e 330, typically associated with rounding in the
 35 instance of a "half-way" input signal value. That is, values for

1 37359/JFO/B600

5 α 310 are generated such that, to the greatest extent possible,
 $E(\alpha) - E(e) = 0$. In addition to choosing values of α intended
to substantially eliminate precision reduction error signal e
330, the values for α also can be selected to tailor the variance
of α to be substantially equal to the variance of the precision
10 reduction error signal, e 330, i.e., $\sigma_\alpha^2 = \sigma_e^2$.

In many cases where the signal values being rounded are
random variables in a binary format, including, without
limitation, fixed point two's complement binary values, $E(e)$ is
15 approximately equal to $\frac{1}{2}$. Thus, it is preferred that α 310 be
assigned values to effectuate an $E(\alpha)$ approximately equal to $\frac{1}{2}$.
One contemplated technique for realizing an $E(\alpha)$ approximately
equal to $\frac{1}{2}$, is to select values for a binary α 310, such that it
20 is substantially equally probable that $\alpha = '1'$ and $\alpha = '0'$. This
can be accomplished using several approaches which utilize
alternating '1' and '0' values for α 310. In one approach, $E(\alpha)$
= $\frac{1}{2}$ can be substantially achieved by assigning to α 310 binary
25 values of '1' and '0' in a strictly alternating, or toggling,
relationship. That is, each successive value of α 310 is either
a '1' followed by a '0', or vice versa. In a second approach,
30 $E(\alpha) = \frac{1}{2}$ can be substantially achieved by assigning to α 310,
binary values of '1' and '0' in an alternating relationship that
employs a selected bit sequence. In a third approach, $E(\alpha) = \frac{1}{2}$
can be substantially achieved by assigning to α 310 binary values
35 of '1' and '0' in an alternating relationship that employs a

1 37359/JFO/B600

pseudorandom bit sequence. Regardless of the approach selected,
a skilled artisan would realize that the commonality among the
5 several approaches is selecting values for α 319 to achieve $E(\alpha)$
 $= \frac{1}{2}$, i.e., that α 310 is assigned both '1' and '0' binary values
with substantially equal probability. The skilled artisan also
10 would understand that, because RTC rounding is commonly used to
effect rounding during the half-value threshold rounding state,
a binary value of '1' is *always* added to the rounding operand,
thus leading to an Expected Value of RTC rounding $E(RTC) = 1$. In
15 the circumstance where $E(e) = \frac{1}{2}$, clearly the commonplace RTC
rounding can account for a significant amount of rounding error,
where SBR rounding generally does not.

Moreover, values for selectable bias α 310 corresponding to
20 an $E(\alpha) = v$, can be chosen to supply a predefined offset value
 v in response to a predetermined characteristic other than the
precision reduction error signal mean value, $E(e)$, such as, for
example, when it is desirable to impart a predefined offset v
25 upon the output signal X 320, or, as another example, when it is
desirable to negate a pre-existing offset having a mean value of
 $-v$ that is present in input signal X 305. Also, values for
selectable bias α 310 can be provided to substantially eliminate
30 precision reduction error signal e 330, and supply a predefined
offset value, v , contemporaneously. That is, selectable bias α
310 can be generated such that $E(\alpha)$ is substantially equal to
35 $E(e) + v$. Furthermore, values for α 310 can be selected to respond

1 37359/JFO/B600

to other predetermined characteristics of X 305, X 320, e 330,
or α 310, as may be indicated by, for example, parametric
5 analysis of the respective signals. A skilled artisan also will
realize that the above techniques are not limited solely to
signals represented by binary values, but can be used, *mutatis*
10 *mutandi*, with values represented in other base systems, whether
or not in complement, and in integer, fixed point, and floating
point formats.

As used in the following examples and Figures, the value of
15 input signal, X , can be represented by a generalized data format
of $s_1.n.a$, where s_1 signifies the sign or polarity of X ; n
signifies an integer part, or digit(s), of the value of X ; "."
signifies the radix point relevant to the value of X ; and a
20 signifies the fractional part, or digit(s), of the value of X .
Similarly, the value of output signal, X , can be represented a
generalized data format of $s_o.n.b$, where s_o signifies the sign or
polarity of X ; n signifies an integer part, or digit(s), of the
25 value of X ; "." signifies the radix point relevant to the value
of X ; and b signifies the fractional part, or digit(s), of the
value of X , with $0 \leq b < a$. A zero-valued a or b indicates that
30 the respective signal has an integer value; a non-zero-valued a
or b indicates that the respective signal is real-valued, with
an integer part and a fractional part.

The following examples and Figures also will be
35 characterized within the context associated with FIG. 2, as

1 37359/JFO/B600

5 further characterized with a fixed point two's complement binary
format. In these examples and Figures, input signal X , with a
precision portion of $(n+a+1)$ bits, provides the rounding operand.
However, because it is desired that the reduced precision output
signal \hat{X} consist of $(n+b+1)$ bits, the most significant $(n+b+1)$
10 bits of the rounding operand are chosen to represent the
precision portion, with the $(n+b+1)$ bit being the least
significant precision digit, and the final $(a-b)$ bits being the
loss portion. In this construct, the rounding point immediately
15 follows the $(n+b+1)$ bit, with the rounding bit being the $(n+b+2)$
digit. When $b = 0$, the rounding point coincides with the radix
point, and \hat{X} is rounded to an integer value. When $b > 0$, the
rounding point follows the b^{th} bit, and the rounding bit is the
20 $b^{\text{th}}+1$ bit. For each signal, both sign bits, respectively s_1 and
 s_0 , a value of binary '0' is indicative of a signal with positive
polarity and a value of binary '1' is indicative of a signal with
negative polarity. In certain embodiments of the present
25 invention, particularly those involving signals in the fixed
point two's complement binary format, and more particularly,
those in which positive and negative values are substantially
equiprobable, it is preferred that the value of input signal sign
30 bit, s_1 , be assigned to selectable bias α when the threshold
rounding state occurs.

FIG. 4 shows an exemplary signal flow model of SBR device
35 400, which receives and combines input signal X 410 with

selectable bias α 420, to produce output signal X 430. The value assigned to selectable bias α 420 is provided by selectable bias generator 440, with the value of α 420 being selected to minimize, or eliminate, the effect of error signal e 450. Without selectable bias generator 440 providing selectable bias α 420, FIG. 4 essentially reduces to model 100 of FIG. 1, in which an error signal e 150 develops as the difference between input signal 110 and output signal 120. To substantially counteract error signal 460 that would otherwise develop between input signal 410 and output signal 430, selectable bias generator 440 assigns values to selectable bias α 420, in response to predetermined characteristics of input signal X 410, output signal X 430, error signal e 460, values of selectable bias α 420, or combinations thereof.

In preferred embodiments of the invention, values of α 420 are chosen such that the Mean, or Expected Value, of α 420, i.e., $E(\alpha)$, is substantially equal to the Mean, or Expected Value, of error signal $E(e)$ 460. Thus, by judiciously selecting values for α , 420 error signal e 460 can be substantially eliminated, i.e., $E(e) - E(\alpha) = 0$. Selectable bias generator 440 also can assign values to α 420 that create offset value v which can, for example, nullify an existing bias in input signal 410, or add a desired offset to output signal 430.

FIG. 5 illustrates method 500, which is another preferred embodiment of the present invention. Initially, the rounding

1 37359/JFO/B600

operand is assigned the value of unassigned input signal X with
5 a digits precision, the operand is defined to have a precision
portion and a loss portion, in a manner consistent with the
reduced precision desired for output signal, \hat{X} with b digits
precision. If $a > b$, then the precision portion of the operand
10 is the first b digits, and the loss portion is the final $a-b$
digits.

At the onset of the rounding operation, the loss portion is
compared with a predetermined threshold value f_t , step 510, which
15 value is representative of a threshold rounding state. Based upon
that comparison, a rounding technique is selected, step 520, to
adjust the value of the least significant precision bit (LSP) of
precision portion 502.

20 If loss portion is substantially equal to predetermined
threshold value f_t , then a selectable bias rounding technique is
used, step 530, in which a selected value is assigned to α ,
responsive to a predetermined characteristic of one of the input
25 signal X ; the output signal \hat{X} ; the error signal e ; the bias α
itself; another preselected characteristic, which may be
determined through parametric analysis of any of these signals;
or a combination thereof. For example, it is be desirable to
30 assign values to α such that $E(\alpha)$ is substantially equal to $E(e)$.
Once bias α is assigned a preselected value, it is combined with
the rounding operand precision portion, step 535.

35

If the loss portion is not substantially equal to the predetermined threshold value f_t , other rounding techniques can be employed, including, for example, RTZ rounding, RTN rounding, RTF rounding, RTC rounding, and RTE rounding, step 540, alone or in combination. At the conclusion of the rounding operation, step 530 or step 540, the rounding operand precision is reduced to $(n+b)$ bits by dropping the final $(a-b)$ loss bits, step 550. The value of the rounding operand is now representative of reduced precision output signal, X , having $(n+b)$ bits precision.

FIG. 6 illustrates another embodiment of the present invention, which describes a method 600 intended to convert real-valued input signal X to a integer-valued reduced precision output value, X . Signal X is represented in $n.a$ format by $(n+a)$ digits, which are assigned to the rounding operand. In this example, it is desired to represent reduced precision output signal X by n digits. Therefore, the most significant n digits of X constitute the precision portion of the rounding operand, with the remaining a digits being the loss portion.

At the onset of the rounding operation, the loss portion is compared with a predetermined threshold value f_t , step 610, which value is representative of a threshold rounding state. In this example, the threshold rounding state is a value half-way between two integers, and f_t is assigned the value of 0.5_{10} . Based upon that comparison, a rounding technique is selected, step 620, to adjust the value of the least significant precision bit.

1 37359/JFO/B600

5 If the loss portion is substantially equal to 0.5_{10} , i.e., f_t , then selectable bias rounding, step 630, is used. In this example, it is preferred that one of two bias values be assigned to α namely, '1' and '0'. Next, the current value of α is combined with the least significant precision digit of the rounding operand to adjust the value of the operand precision
10 portion, step 635.

15 The selectable bias value election is carried out such that α is assigned one value during one pass through step 630, and the other value during an immediately subsequent pass through step 630. The value alternation would continue with each subsequent iteration through step 630, so that the values for and during a series of sequential passes through step 630 would be represented
20 by the sequence {1010...1}. Thus, the rounding operand precision portion is alternately rounded up and rounded down, or toggled, to the nearest integer values. In this case, the expected value of α , $E(\alpha)$, is substantially equal to 0.5_{10} , over a temporal
25 sequence of threshold rounding operations, where the Expected Value of the precision reduction error $E(e)$ also is substantially equal to 0.5_{10} . Thus, the error signal arising from the difference
30 between \hat{X} and X is substantially zero.

Selectable bias rounding step 630 (alternating values) can be implemented in Verilog HDL in a manner similar to true rounding. The Verilog HDL (IEEE Std.1364) is a hardware
35 description language used to design and document electronic

1 37359/JFO/B600

systems, that is well-known to skilled artisans. In this implementation, it is preferred that selectable bias rounding to generate bit sequence for α having substantially zero mean. In one approach, α is inverted every time it is used. The following Verilog HDL code segment models this behavior using a flip-flop that toggles every time α is used.

```
10 wire x[7:0];
    wire x_hat[3:0];
    wire toggle;
15 reg mu;
    assign toggle = (x[3:0] ++ 4'b1000);
    assign x_hat = x[7:4] + (toggle ? mu : x[3]);
    always (posedge clock) mu <= reset_n ? (toggle ? (~mu :
20 mu)) : 1'b0;
```

It is desirable to avoid overflow when performing the addition.

25 The Matlab® integrated technical computing environment can be used to model signal processing methods, and also is well known to skilled artisans. The Matlab® environment is produced by The MathWorks, Natick, MA. Selectable bias rounding of step 730 can be modeled in the Matlab Environment by implementing standard rounding and then performing a correction.

30 The following Matlab® code can implement selectable bias rounding when α toggles between 0 and 1, where \mathbf{x} is a vector of

35

1 37359/JFO/B600

numbers in 1.a format, and \hat{x} is a vector of reduced-precision numbers in 1.b format.

```
5      x_hat = 2^(-b) * floor(x * 2^b + 0.5);
      need_correct = x_hat((x_hat - x) == 2^(-(b+1)));
      mu = zeros(size(correct));
10     mu(2:2:length(mu)) = ones(size(mu(2:2:length(mu))));
      x_hat((x_hat - x) == 2^(-(b+1))) = x_hat((x_hat - x) == 2^(-
(b+1))) - mu * 2^(-b);
```

Instead of using toggled alternating values for selectable bias α , other values and value sequences also can be employed, so long as $E(\alpha)$ resulting therefrom is substantially equal to $E(e)$. One preferred embodiment of an alternating sequence includes the use of a selected sequence of values, such as the exemplary sequence {111000 ... 111}. Another preferred embodiment of an alternating sequence includes the use of a pseudorandom sequence of values, such as the exemplary pseudorandom sequence {0100001001111110}. In either case, the values assigned to α are chosen to produce $E(\alpha)$ substantially equal to $E(e)$.

In a case where it is desirable to impose an offset v upon output signal \hat{x} , values of α may be selected such that $E(\alpha) = 0.5 + v$. Because $E(e) = 0.5_{10}$ in the scenario of method 600, $E(e) - E(\alpha) = v$. For example, if it is desired to add an offset v approximately equal to 0.25_{10} to \hat{x} , then an alternating sequence represented by {1, 0, 0, 0, 1, 0, 0, 0, ..., 1} could be used.

1 37359/JFO/B600

Similarly, a value sequence represented by {0, 1, 1, 1, 0, 1, 1,
1, ..., 0} could be assigned to α to add to X an offset v
5 approximately equal to 0.75_{10} . Furthermore, as before, values can
be assigned to α in any selected sequence, or pseudorandom
sequence, provided the desired values for $E(\alpha)$ are attained.

10 Continuing in method 600, if the loss portion of the
rounding operand is not substantially equal to 0.5_{10} , then
existing rounding techniques can be employed, step 640,
including, for example, RTZ rounding, RTN rounding, RTF rounding,
15 RTC rounding, and RTE rounding, alone or in combination.
Regarding the rounding technique employed at step 640, it is
desirable to choose a technique that does not itself impose an
undesirable error bias upon reduced precision output signal X .
20 At the conclusion of the rounding operation of step 630 or step
640, the rounding operand precision is reduced to n digits by
dropping the least significant a loss digits, step 650. The value
of the rounding operand is now representative of reduced
25 precision output signal, X , having n digits precision.

A skilled artisan would also realize that method 600 is
suitable for implementation in other base environments, including
30 binary. In a corresponding binary implementation, f_t can be
assigned the fractional value of $(.100...0)$, which also is
representative of a half-value state, where $E(e) = \frac{1}{2}$. Thus, it
would be preferred to assign values to the selectable bias such
35 that $E(\alpha)$ is substantially equal to $\frac{1}{2}$. Such values could be

1 37359/JFO/B600

represented by the alternating value sequence represented by {1,
0, 1, 0, ..., 1}.

5 Moreover, in view of the foregoing, a skilled artisan would appreciate that other value sequences can be used in the binary environment for selectable bias α including selected sequences and pseudorandom sequences; and that it also is possible to
10 impose upon reduced precision output signal an offset value such as ∇ , alone and responsive to a predetermined characteristic of, for example, input signal X , output signal \hat{X} , selectable bias α
15 and the like.

FIG. 7 illustrates another embodiment of the present invention, which describes a method 700 intended to convert real-valued, fixed point two's complement binary input signal X to
20 real-valued, fixed point two's complement binary reduced precision output signal \hat{X} . Signal X can be represented by $(n+a+1)$ bits in $s_1n.a$ format, with a fractional part of a bits. Signal \hat{X} is represented by $(n+b+1)$ bits in the $s_1n.b$ format, with a
25 fractional part of b bits, and, $a > b$. To begin, signal X is assigned to the rounding operand. In this example, the most significant $n+b+1$ bits of X constitute the precision portion of the rounding operand. The $n+b+1$ bit is the LSP bit, after which
30 the rounding point is located. The remaining $a-b$ bits are designated as the loss portion of the rounding operand. It may not be necessary, or desired, to include the sign bit in the
35 determination of the number of precision bits and loss bits of

1 37359/JFO/B600

the rounding operand. In this case, the rounding operand would initially be assigned $(n+a)$ bits, of which only $(n+b)$ bits are designated as the precision bits. As before, the least significant $(a-b)$ bits of the rounding operand are designated as the loss bits. In either scenario, sign bit s_1 is not ignored, but is used in the determination of selectable bias α as shown below.

At the onset of the rounding operation, the loss portion is compared with a predetermined threshold value f_t , step 710, which value is representative of a threshold rounding state. In this example, the threshold rounding state is a value half-way between two real-valued binary numbers, and f_t is assigned the value of $(10...0_2)$, i.e, a logical '1' followed by an unbroken series of logical '0'. Based upon the comparison in step 710, a rounding technique is selected, step 720, to adjust the value of the least significant precision bit of the rounding operand.

If the loss portion is substantially equal to $(10...0_2)$, then selectable bias rounding, step 730, is used. In this example, it is preferred that the sign bit s_1 of X be assigned as the bias values. When X is a signal with positive polarity (s_1^+), α is assigned the value of binary '0'; similarly, when X is a signal with negative polarity (s_1^-), α is assigned the value of binary '1'. The technique associated with step 730 can be particularly desirable when X assumes positive and negative values with substantially equal probability. If subsequent values of X are

1 37359/JFO/B600

alternately positive and negative values, then the assignment of
values for α resembles the toggled alternating relationship
5 described relative to FIGS. 5 and 6. In general, however, sign
bit s_1 is a random variable and, thus, provided $E(s_1^+) = E(s_1^-)$,
 $E(\alpha)$ will be substantially equal to $E(e)$. That is, $E(\alpha) = E(e)$
10 $= \frac{1}{2}$. In this case, the assignment of values for α resembles the
use of pseudo random sequence. Once α is assigned the value of
 s_1 , the selectable bias α is combined with LSP bit of the
rounding operand, step 735.

15 On the other hand, if the loss portion of the rounding
operand is not substantially equal to f_t , then existing rounding
techniques can be employed, step 740, including, for example, RTZ
rounding, RTN rounding, RTF rounding, RTC rounding, and RTE
20 rounding, alone or in combination. As before, it is desirable to
choose a rounding technique at step 740, which does not itself
impose an undesirable error bias upon reduced precision output
signal X .

25 At the conclusion of the rounding operation of step 730 or
step 740, the rounding operand precision is reduced to $(n+b+1)$
bits by dropping the least significant $(a-b)$ loss bits, step 750.
30 The value of the rounding operand is now representative of the
signed, real-valued reduced precision output signal, X , having
 $(n+b+1)$ bits precision (or $(n+b)$ bits precision, if the sign bit
is not employed in the characterization of signal "precision.")

35

1 37359/JFO/B600

The following Verilog HDL code segment models the behavior of step 730. If the input X is known to be positive and negative with equal probability, then the sign s_1 of X can be used to create the zero-mean sequence of α as shown below.

```
5      wire x[7:0];  
      wire x_hat[3:0];  
10     wire alpha;  
      assign alpha = x[7];  
      assign x_hat = x[7:4] + ((x[3:0] == 4'b1000) ? alpha :  
15 x[3]);
```

It is desirable to avoid overflow when performing the addition.

A skilled artisan would also realize that method 700 is suitable for implementation in other base environments, including decimal, particularly in complement formats in which the polarity of the values are represented by digits and where positive and negative numbers occur with equiprobability. Moreover, in view of the foregoing, a skilled artisan would appreciate that the foregoing step 730, i.e., assigning the value of s_1 to α can be used in conjunction with the selectable bias techniques described with regard to FIG. 5 and FIG. 6, for imparting a desired offset value ν upon reduced precision output signal X , step 760. The value of ν can be independent of other signal parameters, or can be responsive to a predetermined characteristic of, for example, input signal X , output signal X , selectable bias α and the like.

1 37359/JFO/B600

Furthermore step 780 can be performed sequentially with step 730,
or in tandem therewith.

5

FIG. 8 illustrates another embodiment of the present invention, which describes a method 800 intended to convert a signed, real-valued input signal X to a signed, integer-valued reduced precision output value, \hat{X} . The value of signal X is represented in $s_1n.a$ format by $(n+a+1)$ digits, and is assigned to the rounding operand. In this example, it is desired to represent signal X by n digits. Therefore, the most significant n digits of X , i.e., the integer part, constitute the precision portion of the rounding operand, with the remaining a digits being the loss portion. Digit n is the LSP digit, after which is located the rounding point.

20

At the onset of the rounding operation, the value of the loss portion is compared with a predetermined threshold value f_t , step 810, which value is representative of a threshold rounding state. In this example, the threshold rounding state is chosen to be a value half-way between two integers, (k) and $(k+1)$, and f_t is assigned the value of 0.5_{10} . Based upon that comparison, a rounding technique is selected, step 820, to adjust the value of the least significant precision bit.

30

If the loss portion is substantially equal to 0.5_{10} , i.e., f_t , then selectable bias rounding, step 830, is used. In this example, it is preferred that, if s_1 is indicative of a positive number, then value of the selectable bias α is chosen to round

35

1 37359/JFO/B600

the rounding operand towards negative infinity. On the other hand, if s_1 is indicative of a negative number, then it is preferred that the α is chosen to round the rounding operand towards positive infinity.

As with method 700, method 800 executes the selectable rounding step responsive to sign bit s_1 . Indeed, when the signal values are represented in the two's complement format, the two methods are equivalent in effect because when $s_1 = 1$, i.e., X has negative polarity, a selectable bias value of '1' is combined with the least significant precision bit to adjust the value of the rounding operand, effectively rounding the value towards positive infinity. Similarly, when $s_1 = 0$, i.e., X has positive polarity, a selectable bias value of '0' is combined with the least significant precision bit to adjust the value of the rounding operand, effectively rounding the value towards negative infinity. The current value of α is then combined, step 835, with the least significant precision digit of the rounding operand to adjust the value of the operand precision portion.

As with method 700 at step 730, the technique associated with step 830 can be particularly desirable when X assumes positive and negative values with substantially equal probability. In each method, it is desired to assign values to the selectable bias α so that $E(\alpha)$ is substantially equal to $E(e)$.

35

Continuing in method 800, if the loss portion of the rounding operand is not substantially equal to f_t , i.e., 0.5_{10} , then existing rounding techniques, can be employed, step 840, including, for example, RTZ rounding, RTN rounding, RTF rounding, RTC rounding, and RTE rounding, alone or in combination. Regarding the rounding technique employed at step 840, it may be desirable to choose a technique that does not itself impose an undesirable error bias upon reduced precision output signal \hat{X} .

At the conclusion of the rounding operation of step 830 or step 840, the rounding operand precision is reduced to n digits by dropping the least significant a loss digits, step 850. The value of the rounding operand is now representative of reduced precision output signal, \hat{X} , having s_n digits precision.

FIG. 9 illustrates another embodiment of the present invention, which describes a method 900 intended to convert real-valued fixed point two's complement binary input signal X to real-valued fixed point two's complement binary reduced precision output signal \hat{X} . Signal X is represented by $(n+a+1)$ bits in $s_1n.a$ format, with a fractional part of a bits. Signal \hat{X} is represented by $(n+b+1)$ bits in the $s_1n.b$ format, with a fractional part of b bits, i.e., $a > b$. Signal X is assigned to the rounding operand. In this example, the most significant $(n+b+1)$ bits of X constitute the precision portion of the rounding operand. The $(n+b+1)$ bit is the least significant

1 37359/JFO/B600

precision bit, after which the rounding point is located. The
remaining **(a-b)** bits are designated as the loss portion of the
5 rounding operand.

At the onset of the rounding operation, the loss portion is
compared with a predetermined threshold value f_t , step 910, which
10 value is representative of a threshold rounding state. In this
example, the threshold rounding state is a value half-way between
two real-valued binary numbers, and f_t is assigned the value of
(10...0₂), i.e., a logical '1' followed by an unbroken series of
15 logical '0'. Based upon the comparison in step 910, a rounding
technique is selected, step 920, to adjust the value of the least
significant precision bit of the rounding operand.

If the loss portion is substantially equal to (10...0₂),
20 then in this embodiment of the invention two (or more) selectable
bias rounding techniques may be selected, step 930. In this
example, two values, α_1 and α_4 , may be chosen in the SBR steps
932, 934, for assignment to selectable bias α during successive
25 iterations through step 930, in response to one or more
predetermined characteristics of, for example, input signal X ,
output signal X , the error signal e , selectable bias α , or a
combination thereof. While at least one SBR method of the
30 present invention is preferred to be included in step 930, e.g.,
at step 932, the other rounding technique, used in conjunction
with the SBR method, e.g., step 934, can be a suitable prior art
35 rounding method, if desired. As with methods 500, 600, 700 and

1 37359/JFO/B600

800, α is then combined, step 935, with the least significant precision bit of the rounding operand precision portion.

5 If the loss portion is not substantially equal to $(10...0_2)$, then, unlike methods 500, 600, 700 and 800, step 940 also may employ SBR techniques according to the present invention, also
10 if multiple selectable bias rounding techniques may be selected, step 930, also in response to one or more predetermined characteristics of, for example, input signal X , output signal \hat{X} , the error signal e , selectable bias α , or a combination
15 thereof. In this example, two selectable bias values, α_2 and α_3 , may be chosen in the SBR steps 942, 944, respectively, for the selectable bias α during successive iterations through step 940. While at least one SBR method of the present invention is
20 preferred to be included in step 940, e.g., at step 942, the other rounding technique, e.g., step 944, can be a suitable prior art rounding method. Once a value for bias α is assigned, then it is combined, step 945, with the LSP bit of the rounding
25 operand precision portion, as in step 935. At the conclusion of the rounding operation of step 935 or step 945, the rounding operand precision is reduced to $(n+b+1)$ bits by dropping the least significant $(a-b)$ loss bits, step 950. The value of the
30 rounding operand is now representative of the real-valued reduced precision output signal, \hat{X} , having $(n+b+1)$ bits precision.

FIG. 10 illustrates an selectable bias rounding (SBR) device
35 1000 intended to reduce the precision of input signal 1010, from

1 37359/JFO/B600

data input source 1005. Input signal X 1010 is represented in
an $n+a$ digit format, where n represents the number of integer
5 digits in the value of X and a represents the digits of the
fractional portion 1030 of X . For this example, it is desired
to transform input signal X 1010 into output signal X 1045, which
10 is represented in an $n+b$ digit format, where n represents the
number of integer digits in the value of X and b represents the
digits of the fractional portion of X , where $a > b$. It is also
desired to preset the threshold value f_t within selectable bias
15 generator 1035 to indicate a half-value state, which, in a binary
implementation would be (10....0).

When signal 1010 is admitted to SBR device 1000, fractional
portion a , is compared preselected threshold value f_t within
20 generator 1035, to determine whether a half-value threshold
rounding state exists. If it does, then selectable bias α 1040
is combined with signal 1010 in combiner 1025 to adjust the LSP
of output signal X 1045. Combiner 1025 also drops the least
25 significant $(a-b)$ bits of signal X 1045, which are representative
of the loss portion of the signal. Device 1000 can use input
signal X 1010, selectable bias α 1040, output signal X 1045, or
error signal e 1055, or a combination thereof, to generate values
30 for selectable bias α 1040 that is responsive to a predetermined
characteristic of at least one of the aforementioned signals.
In a preferred embodiment of the invention, values for selectable
35 bias α 1040 are responsive to the error signal e 1055.

FIG. 11 illustrates an embodiment of SBR device 1100 having a particular implementation of bias generator 1135. As in FIG. 10, input signal X 1110 is produced by data source 1105, and received by SBR device 1100. Similar to FIG. 10, input signal X 1110 is represented in an $n+a$ digit format, where n represents the number of integer digits in the value of X , and a represents the digits of the fractional portion 1130 of X . For this example, it is desired to transform input signal X 1110 into output signal X 1145, which is represented in an $n+b$ digit format, where n represents the number of integer digits in the value of X , and b represents the digits of the fractional portion of X , and where $a > b$. It is also desired to preset the threshold value f_t within controller 1137 of selectable bias generator 1135 to indicate a half-value state, which, in a binary implementation would be (10....0).

When signal 1110 is admitted to SBR device 1100, fractional portion a 1130, is compared within controller 1137 to determine whether a half-value threshold rounding state exists. If it does, then controller 1137 enables memory 1138, and a selectable bias α 1140 is produced therefrom, responsive to a predetermined characteristic of one or more of signal X 1110, bias α 1140, signal X 1145, and signal e 1155. Memory 1138 can be, for example, a RAM, a ROM, or a content addressable memory. A skilled artisan will realize that any functionally-equivalent storage device also would be suitable. Memory 1138 can be

1 37359/JFO/B600

5 suitable for use where it is desired to provide particular values for bias α 1140, including selected sequences of alternating values and pseudorandom selectable bias sequences. Selectable bias α 1140, is combined with signal 1110 in combiner 1125 to adjust the least significant precision bit of output signal X 1145. Combiner 1125 also drops the least significant (a-b) bits of signal X 1145, which are representative of the loss portion of the signal. In a preferred embodiment of the invention, values for selectable bias α 1140 are responsive to the error signal e 1155.

15 Figure 12 is yet another embodiment of SBR device 1200 having a particular implementation of bias generator 1235, which is adapted to produce a selectable bias α 1240 having values that alternate, in a toggle relationship, between binary '1' and binary '0'. Input signal X 1210 is produced by data source 1205, and received by SBR device 1200. Similar to FIG. 10 and FIG. 11, input signal X 1210 is represented in an $n+a$ digit format, where n represents the number of integer digits in the value of X and a represents the digits of the fractional portion 1230 of X . For this example, it is desired to transform input signal X 1210 into output signal X 1245, which is represented in an $n+b$ digit format, where n represents the number of integer digits in the value of X and b represents the digits of the fractional portion of X , where $a > b$. It is also desired to preset the threshold value f_t within comparator 1236 of selectable bias generator 1235

1 37359/JFO/B600

to indicate the presence of a half-value state, which, in a binary implementation would be (10....0).

5 When signal 1210 is admitted to SBR device 1200, fractional portion **a**, is evaluated within comparator 1236 to determine whether a half-value threshold rounding state exists. If it does, then comparator 1236 enables flip-flop 1237, and a
10 selectable bias α 1240 is produced therefrom. The present value of selectable bias α 1240 is complemented in inverter 1238 and fed back to the input of flip-flop 1237 to be used as the next
15 value of selectable bias α 1240, when the next threshold rounding state is detected. In this manner, it is possible to implement, for example, the variant of step 630 in method 600 in which a strictly alternating sequence of binary '1' and binary '0' values
20 constitutes successive values for selectable bias α 1240. Selectable bias α 1240, is then combined with signal 1210 in combiner 1225 to adjust the least significant precision bit of output signal **X** 1245. Combiner 1225 also drops the least
25 significant (**a-b**) bits of signal **X** 1245, which are representative of the loss portion of the signal.

Figure 13 is yet another embodiment of an SBR device 1300
30 having a particular implementation of bias generator 1235, which is adapted to produce a selectable bias α 1340 by choosing among different SBR methods, responsive to predetermined signal characteristics. the Input signal **X** 1310 is produced by data
35 source 1305, and received by SBR device 1300. Similar to FIG. 10,

1 37359/JFO/B600

FIG. 11, and FIG. 12, input signal X 1310 is represented in an $n+a$ digit format, where n represents the number of integer digits in the value of X and a represents the digits of the fractional portion 1330 of X . For this example, it is desired to transform input signal X 1310 into output signal X 1345, which is represented in an $n+b$ digit format, where n represents the number of integer digits in the value of X and b represents the digits of the fractional portion of X , where $a > b$. It is also desired to preset the threshold value f_t within comparator 1336 of selectable bias generator 1335 to indicate the presence of a half-value state, which, in a binary implementation would be $(10\dots0)$, or in decimal would be 0.5_{10} .

When signal 1310 is admitted to SBR device 1300, fractional portion a 1330, is evaluated within comparator 1336 to determine whether a half-value threshold rounding state exists. If it does, then comparator 1334 indicates the rounding state to bias control 1337 and enables bias storage 1339 in cooperation with bias control 1337. Responsive to the command from bias control 1337, bias storage assigns to selectable bias α 1340 values that are responsive to predetermined characteristics of input signal X 1310, selectable bias α 1340, output signal X 1345, or error signal e 1355, or a combination thereof. Bias storage 1339 may assign values to selectable bias α 1340, that correspond with multiple SBR techniques, in which case, bias control 1337 also is adapted to assign values for selectable bias α 1340 according

1 37359/JFO/B600

to those multiple SBR techniques. In addition to responding to
the case where $a = f_t$, bias generator 1335 is capable of
5 determining whether $a > f_t$ or $a < f_t$, adopting a separate response
for each case.

Selectable bias α 1340, is then combined with signal 1310
10 in combiner 1325 to adjust the least significant precision bit
of output signal X 1345. Combiner 1325 also drops the least
significant $(a-b)$ bits of signal X 1345, which are representative
of the loss portion of the signal.

15 Figure 14 illustrates yet another embodiment of SBR device
1400 according to the present invention in which bias generator
1445 employs an adaptive parametric analyzer 1436, which may be
used to closely follow real-time statistical parameters of input
20 signal X 1410, output signal X 1445, selectable bias α 1440,
precision reduction error signal e 1455, external controller
1470, and combinations thereof. In this particular embodiment
of the present invention the desired threshold value f_t can be
25 maintained in threshold memory 1437. Once analyzer 1436 has
evaluated one or more of signal X 1410, output signal X 1445,
selectable bias α 1440, and precision reduction error signal e ,
30 it provides a signal to bias selector 1438 which, in turn,
assigns values to selectable bias α 1440 in response to the
analyzed signals. In addition, analyzer 1436, bias selector
1438, or both, may assign values of selectable bias α 1440 at
35 least in partial response to external controller 1470. Bias

1 37359/JFO/B600

selector 1438 may assign values of selectable bias α 1440, that
correspond with multiple SBR techniques responsive to information
5 from parametric analyzer 1436. In addition to responding to the
case where $a = f_t$, bias generator 1435 is capable of determining
whether $a > f_t$ or $a < f_t$, adopting a separate response.

10 FIG. 15 illustrates yet another embodiment of an SBR device
1500 according to the present invention. In this example, input
signal X 1510 is represented in fixed point two's complement
format. Signal X 1510 has a 6-bit fractional portion a_5-a_0 (1511-
15 1516) which, when being rounded to output signal \hat{X} 1545, must be
reduced to a 3-bit fractional portion b_2-b_0 (1546-1548) with
rounding point 1590 conceptually following bit a_3 1513.

Fractional bits a_5-a_0 1511-1516 are compared in comparator
20 1520 with the preselected threshold value f_t 1522, which value is
maintained in storage device 1524. Where it is desired to
implement a sign-bit rounding technique, similar to one discussed
relative to method 800 and FIG. 8, the value of sign bit s_1 1530
25 is selected by MUX 1532 to be combined with the precision portion
1534 of the output of comparator 1520, which constitutes the
rounding operand. If the sign-bit technique is selected by bias
control 1550, then when input signal X 1510 has positive
30 polarity, then the value of binary '0' becomes selectable bias
 α 1533 which is combined with bit a_3 1513. On the other hand,
when input signal X 1510 has negative polarity, then the value
35 of binary '1' becomes selectable bias α 1533 which is combined

1 37359/JFO/B600

with bit a_3 1513. If the current rounding state is not the
threshold rounding state, then bit a_2 1514, can be selectively
5 added to bit a_3 1513 under the direction of bias controller 1550,
in a manner consistent with the rounding technique chosen (e.g.,
RTZ, RTN, RTC, RTF, and RTE). Adder 1535 is designed with a
10 seven-bit input and seven bit output, thus inherently dropping
loss bits a_2 - a_0 (1514-1516), and producing signed output signal
1545 with 3 bit fractional portion b_2 - b_0 1548.

Bias controller 1550 can select from among multiple SBR
15 techniques as well as multiple prior art rounding techniques in
order to effect precision reduction responsive to input signal
1510, output signal 1545, a model of the precision reduction
error as programmed into bias controller 1550, and so forth. For
20 example, where it is desired that bias controller 1550 assign
different values to selectable bias α 1533 depending upon whether
comparator control signal 1570 ("LEG" signal) indicates whether
 $a < f_t$, $a = f_t$, or $a > f_t$, and bias controller 1550 selects the
25 precision reduction techniques for which the controller 1550 has
been programmed.

Figure 16 shows an improved SBR arithmetic unit 1600 which
30 combines an existing arithmetic unit 1605 with a SBR device 1610.
Such an arithmetic unit could be, for example, an arithmetic
logic device 1605 whose output 1607 produces a datum with more
precision, or more digits, than is desired. SBR device 1610 can
35 be used both to reduce the precision of signal 1607, and to

1 37359/JFO/B600

compensate for precision reduction error bias. Also, device 1610
can be employed to compensate for undesirable biases and offsets
5 that may be introduced by data source 1620 and data source 1640.
Furthermore, device 1610 can be adapted to compensate for
computational errors which arise from the operations of
10 arithmetic unit 1605, or to impart a desired offset upon output
signal 1650.

Arithmetic unit 1605 can be, for example, a multiplier, an
adder, an accumulator, or other arithmetic device. In the
15 example shown in Figure 16, arithmetic unit 1605 can be a
multiplier which receives signal 1625, from data source #1 1620,
having p_1 digits, and multiplies signal 1625 with signal 1645
from data source #2 1640, having a precision of p_2 digits.
20 Arithmetic unit 1607 outputs signal 1607 having precision of p_1+p_2
digits. SBR device 1610 can be advantageous where it is desired
to provide data sink 1660 with reduced precision signal 1650, for
example, with p_3 bits precision, where $p_3 < p_2+p_1$.

25 Depending upon the application at hand, a skilled artisan
would be able to implement SBR device 1610, for example, using
one of those illustrated in FIGS. 10-15, to produce the desired
30 results, or another suitable hardware design implementing methods
according to the present invention.

In Figure 17, four-tap LMS adaptive filter 1700 according
to the present invention is illustrated. Overall, filter 1700
35 demonstrates a well-known general architecture for LMS adaptive

1 37359/JFO/B600

filters. Filter 1700 includes FIR filter taps, as illustrated
by filter tap 1718, which includes a number of elements that tend
5 to produce results with increased precision, which may need to
be reduced prior to the datum entering a subsequent processing
element. Such elements can include multipliers 1720, 1730, adder
10 1732, and accumulator 1724. Each of these elements tend to
require greater precision on their respective outputs, and the
judicious placement of SBR devices 1740, 1742, 1744, 1746 in the
data path following selected multipliers 1720, 1730, adder 1732,
15 and accumulator 1724 can be advantageous for managing, or
substantially eliminating, precision reduction errors throughout
filter 1700. Also, SBR devices 1748, 1750, 1752 may be employed
at other points within LMS filter 1700, where it is desirable to
20 substantially eliminate precision reduction errors. Furthermore,
composite elements such as a shifter, multiplier-accumulator
(MAC), or other computational device may benefit from having
certain implementations of the present invention coupled thereto.
25 Finally, it may be desirable to provide SBR arithmetic units
which advantageously combine an existing precision-increasing
arithmetic element with an embodiment of the precision reduction
invention herein.
30

Many alterations and modifications may be made by those
having ordinary skill in the art without departing from the
spirit and scope of the invention. Therefore, it must be
35 understood that the illustrated embodiments have been set forth

1 37359/JFO/B600

only for the purposes of example, and that it should not be taken
as limiting the invention as defined by the following claims. The
5 following claims are, therefore, to be read to include not only
the combination of elements which are literally set forth but all
equivalent elements for performing substantially the same
10 function in substantially the same way to obtain substantially
the same result. The claims are thus to be understood to include
what is specifically illustrated and described above, what is
conceptually equivalent, and also what incorporates the essential
15 idea of the invention.

20

25

30

35